

AD-A147 520


PRINCIPLES OF MULTIOBJECTIVE OPTIMIZATION(U) NAVAL
POSTGRADUATE SCHOOL MONTEREY CA R E ROSENTHAL AUG 84
NP555-84-016


1/1

UNCLASSIFIED

F/G 5/1

NL





AD-A147 520

2

NPS55-84-016

NAVAL POSTGRADUATE SCHOOL

Monterey, California



DTIC
ELECTE
OCT 31 1984
B

PRINCIPLES OF MULTIOBJECTIVE OPTIMIZATION

by

Richard E. Rosenthal

August 1984

Approved for public release; distribution unlimited

Prepared for:
Naval Postgraduate School
Monterey, California 93943

DTIC FILE COPY

84 10 30 096

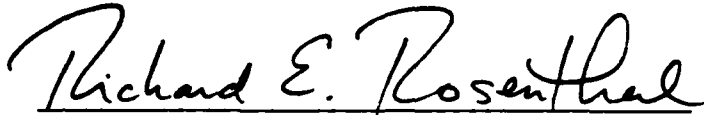
NAVAL POSTGRADUATE SCHOOL
Monterey, California

Commodore R. H. Shumaker
Superintendent

David A. Schradly
Provost


This work was supported by the Naval Postgraduate School Foundation
Research Program under contract with the National Research Council.

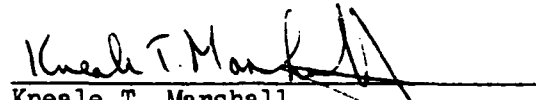
Reproduction of all or part of this report is authorized.


Richard E. Rosenthal, Senior NRC Associate
Department of Operations Research

Reviewed by:

Released by:


Alan R. Washburn, Chairman
Department of Operations Research


Kneale T. Marshall
Dean of Information and Policy

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER NPS55-84-016	2. GOVT ACCESSION NO. AD-A147520	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) PRINCIPLES OF MULTIOBJECTIVE OPTIMIZATION		5. TYPE OF REPORT & PERIOD COVERED Technical
7. AUTHOR(s) Richard E. Rosenthal		6. PERFORMING ORG. REPORT NUMBER
9. PERFORMING ORGANIZATION NAME AND ADDRESS Naval Postgraduate School Monterey, CA 93943		8. CONTRACT OR GRANT NUMBER(s)
11. CONTROLLING OFFICE NAME AND ADDRESS		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 61152N; RR000-01-10 N0001484WR41001
13. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		12. REPORT DATE August 1984
		13. NUMBER OF PAGES 36
		14. SECURITY CLASS. (of this report) UNCLASSIFIED
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Optimization Multiple Objectives (Multicriterion Decision Making) in management type of research.		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) The purpose of this paper is to isolate and analyze the principal ideas of multiobjective optimization. This is done without casting recriminatory aspersions on single-objective optimization or championing any one multiobjective technique. The paper first attempts to define the problem, and then discusses the fundamental ideas, many of which stem from common sense. Each idea is examined for strengths and weaknesses and two—efficiency and utility—are shown worthy of extended consideration. In the light of this analysis some general recommendations are made. Besides offering the simple advice of not dismissing		

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

single-objective optimization as a possible approach, we suggest that three broad classes of multiobjective techniques are very promising in terms of reliably and believably achieving a most preferred solution. These are: (i) partial generation of the efficient set, a rubric we use for unifying a wide spectrum of existing methods; (ii) explicit utility maximization, a much overlooked approach combining multiattribute decision theory and mathematical programming; and (iii) implicit utility maximization, the name we use for the popular class of methods introduced by Geoffrion, Dyer and Feinberg and extended significantly by many others.

3

PRINCIPLES OF MULTIOBJECTIVE OPTIMIZATION

Richard E. Rosenthal
National Research Council Senior Research Associate*
Operations Research Department
Naval Postgraduate School
Monterey, California 93943

*On leave from position as Associate Professor of Management Science
at the University of Tennessee.

AUGUST
~~July~~ 1984

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
PER CALL JC	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

PRINCIPLES OF MULTIOBJECTIVE OPTIMIZATION

Richard E. Rosenthal

July 1984

The purpose of this paper is to isolate and analyze the principal ideas of multiobjective optimization. This is done without casting recriminatory aspersions on single-objective optimization or championing any one multiobjective technique. The paper first attempts to define the problem, and then discusses the fundamental ideas, many of which stem from common sense. Each idea is examined for strengths and weaknesses and two--efficiency and utility--are shown worthy of extended consideration. In the light of this analysis some general recommendations are made. Besides offering the simple advice of not dismissing single-objective optimization as a possible approach, we suggest that three broad classes of multiobjective techniques are very promising in terms of reliably and believably achieving a most preferred solution. These are: (i) partial generation of the efficient set, a rubric we use for unifying a wide spectrum of existing methods, (ii) explicit utility maximization, a much overlooked approach combining multiattribute decision theory and mathematical programming, and (iii) implicit utility maximization, the name we use for the popular class of methods introduced by Geoffrion, Dyer and Feinberg and extended significantly by many others.

1. Introduction

A common problem faced by decision makers is choosing the best or optimal course of action among feasible alternatives. However reasonable this desire may seem, there can be disagreement in some cases as to what it means. The basic problem is in defining the "best" solution.

This is not always a difficulty. There have been many situations in which the best course of action is well-defined: namely, the alternative that maximizes or minimizes a well-defined (scalar-valued) objective function. Such cases comprise the mainstream of management science/operations research applications.

But in other decision making problems, there is no single objective function which can adequately serve to compare the difference in desirability between feasible solutions. Multiobjective optimization is a subfield of management science/operations research that deals with these problems.

There are several reviews of multiobjective techniques available (e.g., Cohon [1978], Cohon and Marks [1975], Evans [1983], Goicoechea et al. [1982], Haimes et al. [1975], Harrison [1983], Ho [1979], Hwang et al. [1979], Roy and Vinke [1981], and Zeleny [1982]). Various methods for taxonomizing the literature have been proposed. For example, one way to classify techniques is according to whether the decision maker interacts with the solution process before, during or after the analysis (or not at all). Another useful classification by Ho is based on the types of information required of the decision maker. The purpose of this paper is not to survey techniques but rather to isolate and analyze the principal ideas upon which the techniques are based. The conclusions drawn at the end of this paper are recommendations of broad classes of multiobjective methods, not of specific techniques. We believe it is very important to approach this area without an overstated opinion of its power relative to single-objective optimization, and without a strong feeling of advocacy towards any one multiobjective technique.

1.1 Terminology and Scope

The term multiobjective optimization is not used universally in the literature. Other terms such as multiple criteria decision making are more wide-spread, e.g. Zeleny [1982]. Regardless of terminology, what is common to the problems of this field is the existence of multiple

measures (whether called criteria, attributes, objectives or objective functions) of the quality of feasible alternatives.

An attempt to clarify terminology would be useful here. We shall use the term multiple criterion decision making to broadly describe all quantitative decision problems in which multiple measures of solution quality exist. Then we divide multiple criterion decision making into two parts: 1) multiobjective optimization, which refers to problems with a large number of feasible alternatives, and 2) multiattribute decision theory, which refers to problems with a small number of feasible alternatives. We distinguish a large set from a small set operationally: the feasible region is large if the analyst considers it too big to totally enumerate.

A reason for making this differentiation in terminology is that it enables multiobjective optimization to be viewed as an extension of optimization. The shared property of the two is the need to use efficient algorithms for narrowing down a large (possibly infinite) set of feasible alternatives to a small (possibly singleton) set of recommended solutions.

The central scope of this paper is multiobjective optimization, but not because of a lack of respect for multiattribute decision theory. Techniques in the latter area have been widely applied. They normally call for the evaluation of each alternative with respect to each objective at the outset of the analysis. Typical situations amenable to this requirement are problems involving the siting of power plants or other large, complex facilities (e.g., Hobbs [1979], Keeney [1979, 1980]) or the selection of expensive, complex equipment (e.g., Hannan et al. [1983]). Usually only a small number of alternatives (say, 5 to 20) pass the initial economic, environmental and technological screenings of

potential candidates. Methods calling for the evaluation of all feasible alternatives are, by definition, not directly applicable to multiobjective optimization. However, one theme to be developed later in this paper is that multiattribute decision theory is in fact far more applicable to multiobjective optimization than is widely appreciated.

The scope of this paper does not include problems with multiple decision makers. The added difficulties in those problems may be conflicts not only on the importance of the various objectives but also on how to measure them.

Another limitation on the scope of this paper is that we do not directly confront uncertainty. Generally speaking, this issue is treated in greater depth in the multiattribute decision theory literature (e.g., Dyer and Sarin [1979], Farquhar [1983, 1984], Fishburn [1983], Keeney and Raiffa [1976], and Winkler [1982]) than in the multiobjective optimization literature.

1.2 An Attempt at Definition

The mathematical entities required to define an instance of multiobjective optimization are a set of feasible decisions X (contained in R^n) and a set of scalar-valued objective functions, f_i , $i=1, \dots, k$, defined on X . The notation $F(x)$ will be used for the vector function $(f_1(x), \dots, f_k(x))$.

There is no precise mathematical statement for the multiobjective optimization problem, although language like "maximize $F(x)$ over all $x \in X$ " has been used to describe it in the past. Unfortunately, the phrase "maximize $F(x)$ " lacks meaning, because the set $\{F(x): x \in X\}$ lacks a natural ordering whenever F is vector-valued. With no such ordering,

given two feasible alternatives y and z , there may be no definite answer as to whether $F(y)$ is "greater than," "less than," or "equal to" $F(z)$. The inability to specify which of two alternatives is greater implies the inability to specify which of several is greatest.

A meaningful definition of multiobjective optimization requires the use of some subjectivity, regardless of how precisely $F(x)$ and the feasible region X may be defined. A reasonable statement of the problem is: find a feasible x so that the most preferred vector of objective function values $F(x)$ is attained. Here the subjectivity is captured in the term "preferred," which has no rigorous definition. The lack of a rigorous, nonsubjective definition is important to point out because some of the multiobjective optimization literature treats the problem as if it were amenable to cut-and-dried solutions, which it is not. There is in fact no absolute meaning of "best decision" in the multiobjective context.

While the previous point is responsible for attracting psychologists, philosophers and social scientists to multiobjective optimization, one can assert for balance that some of the work in the area could pass mathematical muster with the "purists" of management science/operations research. To achieve a rapprochement between the behaviorally oriented and mathematically oriented camps would in fact be worthwhile. All too often the former group unfairly dismisses single-objective optimization as unrealistic; or, equally inappropriately, they regard single-objective problems as relatively unimportant special cases of multiobjective problems. The record of successful applications of single-objective optimization is far too impressive, however, for anyone to dismiss (or co-opt). On the other hand, those who scoff at multiobjective optimization as too "soft" cannot reasonably ignore the problem nor should they

doubt that useful mathematical analysis can be brought to bear on the problem.

2. Common Sense Approaches

Many of the ingredients of existing techniques for solving multi-objective optimization problems are common sense ideas. These are given below.

2.1 Weights

An obvious approach to multiobjective optimization is to specify a "weight" w_i , representing the importance of the i^{th} objective function, and then to solve

$$\text{maximize } \sum_i w_i f_i(x) \text{ s.t. } x \in X$$

by means of a single-objective optimization algorithm. Though used frequently, this approach has some serious theoretical objections (discussed in Section 4.3). Weighting is still a very important idea and it can be of value when used in other ways. The essential point of caution about weights is that they should not be regarded as constant: the decision maker's perception of the importance of f_i should be allowed to vary as the problem-solving procedure evolves. There exist many techniques for choosing weights; see Hobbs [1979, 1980] and Schoemaker and Waid [1982] for reviews and comparisons.

2.2 Targets, Ideals, Aspiration Levels etc.

Another common sense approach is to set target values for all (or all but one) of the objective functions. Then the achievement or excess of each target is stipulated as a constraint in the ordinary mathematical programming sense; e.g., $f_i(x) \geq b_i$.

Targets can be determined in many ways. One extreme is to set them as the minimal values with which the decision maker might feel satisfied. The opposite extreme is to use each objective's ideal, which is its best attainable value when all other objectives are ignored. Of course, in any realistic setting, there are no feasible solutions which simultaneously attain all the ideals. (This is likely to be true of other sets of targets as well.) Two approaches often taken are either to "displace the ideal" (Zeleny), i.e., aim for lower targets; or to define penalties for failing to meet targets and then minimize the sum of penalties.

In some models targets are expressed as intervals rather than points. This has the advantage of allowing the decision maker to express his or her desires with less precision than a point target.

Target setting is usually an iterative process with adjustments made up or down depending on the under- or over-achievements of a previous iteration. Soland [1979] pointed out the connection between this process and the theory of satisficing due to Simon [1957, 1964].

One note of caution in using ideals, which is perhaps infrequently mentioned, is that they can fail to exist finitely. This is not a matter of splitting mathematical hairs because there exist reasonable cases where this situation arises. Suppose the ideal of f_i is defined as its maximum over X but f_i is unbounded from above on X . In single-objective optimization, this condition normally signals a modeling error, e.g., a constraint omitted or a cost too low. However, in multi-objective optimization the condition is not unlikely in a correct formulation. The aspect of reality which would prevent infinite attainment of f_i may very well be consideration of some conflicting objectives. But these are ignored in the definition of the ideal. It is not clear

whether or not the methods that call for computation of ideals can endure this difficulty. Substituting a large finite number in place of an infinite ideal may not be enough to keep the method on firm footing.

Targets are sometimes used in conjunction with weights, e.g., see the discussion of goal programming in Section 4.3. Hobbs [1979] warned that the targets in models of this type can have consequences not intended by the decision maker, in terms of causing tradeoffs to be made which are not implied by the weights.

2.3 Priorities

Another natural idea for the handling of multiple objectives is to assign priorities to the objectives, and then optimize them one at a time in priority order. At each successive optimization, higher priority objective functions are constrained to their previously optimized level. The technique of pre-emptive goal programming (also known as lexicographic goal programming) is a popular example of this approach.

The idea of this approach is to give the decision-maker the best attainable value of the first priority objective. If the best value is achieved at a unique point in X , the process ends. Otherwise, the tie is broken by choosing a point among the alternative optima that is best with respect to the second priority objective. If ties still exist, the third priority objective is used, and so on.

There are two evident limitations of this way of using priorities. First of all, the successive tie-breaking approach is applicable only when a massive number of ties occur (i.e., dual degeneracy in the mathematical programming context). There may be little reason to expect the process to continue much beyond the first maximization.

Thus, some, if not most, of the objective functions will be totally left out of consideration, a serious shortcoming.

A second weakness in the use of priorities as above is that it disallows the very reasonable practice of trading off a small degradation in a high priority objective for a large improvement in a low priority objective.

The ideas of priorities and targets can be used in conjunction with each other. In methods of this type (e.g., Gilbert and Shane [1981]); constraints on objectives are defined in terms of targets, and priorities are given to these constraints. One then solves a sequence of feasibility-seeking problems, each trying to enforce one additional constraint in priority order. By going through this procedure several times and adjusting the targets sensibly, the weaknesses of prioritization given above can be overcome to a large degree.

2.4 Efficiency

To speak of another common sense concept, efficiency (also known as noninferiority, nondominance, Pareto-admissability and Pareto-optimality), requires an assumption called monotonicity. This means that for every objective function f_i (possibly redefined as $-f_i$), it is assumed the decision maker's satisfaction will never decrease as f_i increases (all other objectives held fixed). This is a reasonable assumption in most instances. An exception, noted by Soland [1979], might arise if f_i were a sugar content objective in a food marketing model. Another exception is the deer population objective in the forest management model of Harrison and Rosenthal [1981]. More deer are normally preferred to fewer for aesthetic and recreational reasons, but not when their popula-

tion is large enough to remove all the forest undergrowth. Despite these instances, monotonicity is assumed throughout the discussions of efficiency to follow.

The definition of efficiency first requires the definition of dominance: given feasible points $y, z \in X$, z is dominated by y if and only if

$$f_i(y) \geq f_i(z), \quad i=1, \dots, k$$

with strict inequality for at least one objective. Then, a feasible solution $x \in X$ is efficient if and only if there does not exist a solution $y \in X$ which dominates x .

Clearly, under the monotonicity assumption, a rational person would never deliberately select a dominated point. This is probably the only important statement in multiobjective optimization that can be made without the possibility of generating some disagreement.

3. More on Efficiency

3.1 Proper Efficiency

If the allegedly noncontroversial declaration "never choose dominated points" were worded "always choose efficient points," then it could in fact be subjected to criticism for not being precise enough. This is due to some interesting analysis of Geoffrion's [1968]. He noted that it would be unwise to select some efficient points, namely those which are not properly efficient. Before we give his definition of proper efficiency, first note, that if $x \in X$ is efficient and some other point y surpasses x in one objective, say $f_i(y) > f_i(x)$, then there must exist some other objective, say f_j , in which x surpasses y .

(Otherwise, x would be dominated by y and could not be efficient.) Let x , y , i and j be defined as in the last statement. Then x is properly efficient if there exists an $M > 0$ such that for any y , i , j so defined

$$\frac{f_i(y) - f_i(x)}{f_j(x) - f_j(y)} \leq M.$$

The ratio in this definition is the improvement in the i^{th} objective divided by the decrement in the j^{th} objective resulting from a change of solution from x to y . If this ratio were not bounded, then an extremely large improvement in f_i could be obtained in exchange for an infinitesimal decrement in f_j . As such an exchange would never be refused by a rational decision maker, one should always choose properly efficient points.

It turns out that the practical significance of the distinction between proper and improper efficiency is not great, but it is a valuable lesson to have analysis rightfully challenge so common-sensical a notion as efficiency. The distinction is actually moot in the important special case when the objectives are linear and X is defined by linear constraints, because then all efficient points are properly efficient (Benson and Morin [1977]). According to Soland [1979], there are other important cases as well in which improperly efficient points are nonexistent or not likely to be considered by the decision maker. (See Benson [1978], however, for a fascinating example in which there are an infinite number of improperly efficient points, yet no properly efficient points.) It is hard to guarantee proper efficiency algorithmically, so Geoffrion recommends that the analyst unsure about the proper efficiency of a selected point should do "stability analysis" (testing of numerous values of the ratio) about the selected point.

3.2 Characterization of Efficiency

There have been numerous researchers including Kuhn and Tucker [1950] and Geoffrion, who discovered a connection between efficiency and the idea of weights given in Section 2.1. They have found, under suitable conditions on the f_i and X , that efficient points correspond to solutions of the scalar optimization problem

$$\max \sum_i w_i f_i(x) \text{ s.t. } x \in X$$

with all $w_i > 0$. An elegant proof based on linear programming duality is given by Isermann [1974] for the all linear special case.

Soland [1979] has given a very general characterization of efficiency in terms of the solutions to scalar optimization problems. Soland's theorem uses the following constructs: let g be any function from R^k to R which is strictly increasing with respect to each component of its argument taken separately (consistent with monotonicity), let $b \in R^k$ (this vector will serve as the targets or aspiration levels as in Section 2.2); and let $P(g,b)$ be the scalar optimization problem:

$$\begin{aligned} &\text{maximize } g(F(x)) \\ &\text{s.t. } x \in X \\ &F(x) \geq b. \end{aligned}$$

In a cunningly simple proof, Soland shows that if x_0 is an optimal solution of $P(g,b)$ then x_0 is efficient; and, conversely, if x_0 is efficient there exists a $P(g,b)$ in which x_0 is optimal. Soland's result does not assume any restrictions such as linearity, convexity, continuity or compactness on the f_i or X . A form of g often used in applications of the theorem is $g(F) = \sum_i w_i f_i$ where all $w_i > 0$. A choice sometimes taken for b_i is $-\infty$.

3.3 Generation of Efficient Solutions

Some authors have suggested that the job of the analyst should be to generate the set of all efficient solutions and to let the decision maker then choose an alternative from the efficient set. This removes the subjectivity and hence any causes for controversy, as the problem of generating the efficient set is well-defined mathematically. Indeed, some authors might argue that the terminology "multiobjective optimization" should be reserved exclusively for this rigorously defined problem. That reservation is not accepted here, however, because just computing the efficient set does not necessarily enable the decision maker to decide on a course of action. Excluding the dominated points is certainly valuable, but the remaining efficient set may be so large that the decision maker is left with a selection problem that is still too difficult to solve unaided.

The case of two objectives is special: the image under F of the efficient set can usually be represented by a curve in R^2 and is thus easy to use. When appropriate convexity conditions hold, it can be generated by solving the parametric program

$$\max w_1 f_1(x) + (1-w_1)f_2(x), \quad \text{s.t. } x \in X.$$

for $0 \leq w_1 \leq 1$. Or in any case it can be generated by solving the parametric program

$$\max f_1(x) \text{ s.t. } f_2(x) \geq b_2, \quad x \in X$$

for $-\infty < b_2 < \infty$, as shown, e.g., by Benson [1979].

Some very clever algorithms have been developed for generating the efficient set in the more difficult case when $k > 2$. Most of this work applies only to the all linear case, e.g., Ecker et al. [1975, 1978, 1980], Gal [1977], Yu and Zeleny [1975]. Even in this special case the structure of the efficient set can be complicated. For example, the points on an edge connecting two efficient extreme points in a linear problem may not be efficient (Ecker and Kouada [1978]); hence, the efficient set is possibly not convex. Ecker and Kouada have devised a simple test for determining whether or not any edge incident to an efficient extreme point is efficient. This test forms the basis of the Ecker-Hegner-Kouada [1980] algorithm for generating maximal efficient faces. A maximal efficient face is a face of the linear polytype X which is efficient and which is not contained within an efficient face of higher dimension. The set of all such faces constitutes a minimal representation of the efficient set.

3.4 Partial Generation of the Efficient Set

Considering the efficient set may be too large to generate (as well as too large to have any managerial use), a number of authors have decided that the best approach to some multiobjective optimization problems is to generate a subset of the efficient set. Included among these are Ecker and Shoemaker [1980]; Gilbert, Holmes and Rosenthal [1982]; Ho [1979]; Hultz, Klingman, Ross and Soland [1981]; Soland [1979]; Steuer and Choo [1981]; and Steuer and Schuler [1978].

We call this approach partial generation of the efficient set and consider it to be among the most promising approaches to multiobjective optimization. There is a wide range of reasonable but differing points

of view as to what the analyst's aim should be in partial generation of the efficient set. At one end of the spectrum is Soland's and Hultz et al.'s' view that the analyst should put hardly any guidelines or restrictions on the decision maker. Rather, the analyst should provide an interactive instrument with which the decision maker can probe, sample and wander through the efficient set at his or her own discretion. The Hultz et al. paper treats a multiobjective facility location problem by simply requiring the decision maker to sit at a terminal and define a series of weights w_i and targets b_i . Each of these sets of information defines a $P(g,b)$ subproblem with linear g . The subproblem is solved by an integer programming algorithm which yields (by Soland's theorem) an efficient solution. The decision maker is then invited to alter the weights or targets and proceed to the solution of another $P(g,b)$. An elegant special-purpose language was designed for carrying out this interaction between the decision maker and the subproblem solver.

The opposite end of the spectrum from Soland's highly unstructured, totally interactive approach is Ecker and Shoemaker's [1981] highly structured, totally automated approach to partial generation. Ecker and Shoemaker's approach is to narrow down the efficient set to a preferred subset analytically. To this end they define types of efficient points that are more desirable than others and then devise algorithms for generating these special subsets. To date, these algorithms apply only to problems with linear objective functions and linear constraints.

The first special type of efficient point that Ecker and Shoemaker consider is the efficient compromise. Let

$$d_i(x) = M_i - f_i(x)$$

where M_i is the ideal of f_i . (The M_i are assumed finite here but recall from Section 2.2 that this may sometimes be difficult to guarantee.) A point $y \in X$ is a compromise solution if it minimizes the maximum deviation from the ideal, i.e., y is an optimal solution in

$$\min_{x \in X} \max_i d_i(x).$$

In general, not all compromise solutions are efficient and vice versa. The set of points with both properties is called the efficient compromise set.

Ecker and Shoemaker's second analytic approach to partial generation of the efficient set is the tradeoff compromise set. This is a subset of efficient points with the additional property: $y \in X$ is a tradeoff compromise point if $f_j(x) > f_j(y)$ for some $x \in X$ implies $f_i(y) > f_i(x)$ and $d_i(y) \geq d_j(y)$ for some i . The idea of the definition is that departing from a tradeoff compromise point to gain an improvement in f_j necessitates a degradation in some f_i that is worse (in terms of deviation from the ideal) than the sacrifice in f_j of not departing. Ecker and Shoemaker show how to generate the tradeoff compromise set for linear multiobjective problems, and they discuss other analytic partial generation methods as well.

A third philosophy of partial generation that lies between the highly unstructured, totally interactive Soland approach and the highly structured, totally analytical Ecker-Shoemaker approach is that of Steuer and Choo [1983] and Steuer and Schuler [1978]. Their idea is to generate a subset of efficient points that are "dispersed" and hence "representative" of the entire set, but they use guidance from the

decision maker as to the "type" of point to aim for. Applying only to the linear multiobjective case, their algorithms involve structured interaction between the decision maker and a parametric programming routine. The holistic preference evaluation method of Ho [1979] is another interactive linear multiobjective technique that generates part of the efficient set through structured dialogue and parametric programming.

In the author's opinion the wide spectrum of techniques for partial generation of the efficient set, such as those mentioned here, are among the most promising methods for dealing with multiobjective optimization problems. A method of Gilbert, Holmes and Rosenthal [1982], which was applied to an integer programming land allocation model, is a recent addition to this class.

4. Utility

In the preceding discussions efficiency emerged as a very important idea which forms the backbone of a variety of multiobjective methods. The next major idea we cover is the utility function. A utility function $U: R^k \rightarrow R$ has the following significance: given $y, z \in X$, we have $U(F(y)) > U(F(z))$ if and only if $F(y)$ is preferred to $F(z)$. Assuming such a function U exists, the multiobjective optimization problem can then be reduced to a single-objective optimization problem:

$$\max U(F(x)) \quad \text{s.t. } x \in X.$$

Techniques that make use of this concept are of two types: explicit utility maximization and implicit utility maximization. (Harrison [1983] has proposed an implicit/explicit utility function approach which

is a hybrid of the two.) The roots of utility-based methods lie firmly in the area of economics. We first explore this connection before describing the explicit and implicit approaches.

4.1 The Relationship Between Multiobjective Optimization and Economic Theory of Demand

A close relationship between the ideas of optimization theory and economic equilibria has been observed since the early years of mathematical programming (e.g., Baumol [1977], Kuhn and Tucker [1950]). The development of some thoughts along these lines is helpful in that multiobjective optimization can then be viewed as a generalization of an economic equilibrium concept.

A classical economic problem (e.g., Baumol) is to assume a consumer has I dollars to spend and must allocate these funds among k commodities. The consumer's individual purchasing power is too small to affect prices so P_i is the price per unit of the i^{th} commodity, which holds regardless of the number of units (a_i) consumed. The problem is to predict the consumption levels a_1, \dots, a_k .

The optimization problem implicit in this prediction is

$$\max U(a_1, \dots, a_k)$$

$$\text{s.t. } \sum_i P_i a_i = I$$

$$a_i \geq 0$$

where U is the utility function representing the satisfaction obtained from a particular commodity combination $a = (a_1, \dots, a_k)$. To proceed with this development it is not necessary to assume knowledge of an

explicit representation of U . It is only necessary to accept some fairly reasonable assumptions about consumer behavior (see Baumol), the results of which are that U is strictly concave.

The solution a^* to this one-row nonlinear program is characterized by the necessary and sufficient optimality conditions

$$MRS_{ij}(a^*) = - P_i/P_j$$

where MRS_{ij} , the marginal rate of substitution, is defined as

$$MRS_{ij}(a) = \frac{\partial U}{\partial a_i} / \frac{\partial U}{\partial a_j}.$$

This standard result is easily derived with a Lagrange multiplier (and can be clearly demonstrated geometrically for the case of $k=2$).

There are two reasons why this simple classical result is of interest here. The first is that the economists are to be credited with a very useful idea: they show us how to recognize an optimal solution for a problem whose objective function we may not know how to specify. What has to be assumed, in addition to concavity, is that the MRS can be evaluated for a specific a . This turns out to be a much smaller information burden on the analyst than deriving an equation for U , because $MRS_{ij}(a)$ is vividly interpretable as the rate at which the consumer would be willing to trade off commodity i for commodity j starting from point a . As discussed later (Section 4.4), Geoffrion, Dyer and Feinberg [1972] were the first to recognize that this idea can be extended to be the basis of an algorithm for multiobjective optimization.

The second reason for present interest in the classical economic theory is that it provides a framework for conceptualizing the multi-

objective optimization. The decision maker in the multiobjective optimization problem is just like the consumer deciding how to fill his or her market basket with commodities, except that now the "commodities" are functions ($a_i = f_i(x)$) and the single explicit constraint on total purchases ($\sum_i p_i a_i = I$) is replaced by a more complicated, composite constraint: $a \in \{F(x): x \in X\}$.

4.2 Explicit Utility Functions

The explicit utility function approach is to assess an explicit form of U by techniques of multiattribute decision theory (e.g., Dyer and Sarin [1979], Farquhar [1984], Fishburn [1983], Keeney [1977], Keeney and Raiffa [1976], Kirkwood and Sarin [1980]) and then to solve

$$\max U(F(x)) \quad \text{s.t. } x \in X$$

directly. The great advantage of this approach is that it makes the vast body of theory, algorithms, software and experience that currently exist for single-objective optimization immediately available for solving multiobjective problems. This is no minor feat when one considers the extent to which (single-objective) optimization has influenced the evolution of computers and computer science.

In spite of this great advantage, the explicit approach has been used very rarely for multiobjective optimization. In contrast, explicit utility functions have been used frequently in multiattribute decision theory. In other words, the explicit approach has been restricted in practice almost entirely to situations where the feasible region is small enough to totally enumerate. This restriction is not due to theoretical limitations of the explicit approach. According to Hobbs

and Keeney [personal communications, 1982], nearly all multiattribute utility assessment techniques can be readily adapted for use in the large-feasible-set case.

The explanation as to why the explicit approach has been taken so infrequently in the large-feasible-set case is probably circumstantial. The fields of multiattribute decision theory and multiobjective mathematical programming have evolved with hardly any interaction in spite of their common aims. We believe that combining the two fields, in the form of the explicit utility function approach, is a promising yet much overlooked approach to multiobjective problems. (See also DeWispelare and Sage [1981].) Those who may have considered and rejected this approach perhaps overestimated the difficulty of assessing an appropriate U. Or perhaps they concluded that mathematical programming algorithms were unsuitable for globally optimizing the common multiattribute utility functions. Some recent analysis of Harrison's [1983] indicates that these functions do, in fact, very frequently possess the desirable feature of strict quasiconcavity, which guarantees that a local optimum is global.

The only instances we know of, in which the explicit utility approach was taken in a multiobjective optimization problem, are reported by Golabi, Kirkwood and Sicherman [1981], Gros [1975], Harrison and Rosenthal [1981], Keefer [1978] and Ringuest and Gullledge [1983]. Golabi et al.'s application was to select a portfolio of solar energy projects. Gros considered power plant siting. Harrison and Rosenthal's application of the idea is in a forest management model which has been used by over 1500 landowners throughout the southeastern United States. Keefer's paper involved industrial resource allocation and Ringuest and Gullledge's

paper described a hypothetical manufacturing example. In all these cases mathematical programming algorithms were used to find a most preferred solution once U was assessed. The form of U chosen was either of the explicit utility functions that Keeney and Raiffa refer to as the additive and multiplicative forms.

The additive and multiplicative forms both require, for each f_i , the assessment of a single-attribute utility function $u_i(f_i)$ which maps achievement of f_i on to the interval $[0,1]$. The value $u_i = 1$ is assigned to the most desirable level of f_i (the ideal), while the value $u_i = 0$ is assigned to the least desirable level. Intermediate values of u_i naturally indicate proportionate intermediate levels of desirability. When monotonicity holds, an exponential form

$$u_i(f_i) = \alpha_i(1 - \exp[\beta_i(f_i - \gamma_i)])$$

is often used for the single-attribute utility function (e.g., Keeney [1979]). The parameters α_i , β_i , γ_i can be computed after assessing $u_i^{-1}(0)$, $u_i^{-1}(0.5)$ and $u_i^{-1}(1)$. (The decision theoretic principle underlying the exponential form is "constant risk preference;" see Keeney and Raiffa [1976]).

The additive form of $U(F)$ requires the assessment of weights and is simply

$$U(F) = \sum_i w_i u_i(f_i)$$

(not to be confused with $\sum_i w_i f_i$ as in the weighting method). The multiplicative form of the utility function requires parameters λ , $\lambda_1, \dots, \lambda_k$ in addition to the single-attribute functions. It takes the form

$$U(F) = \frac{1}{\lambda} [-1 + \prod_i (1 + \lambda \lambda_i u_i(f_i))]]$$

where

$$1 + \lambda = \prod_i (1 + \lambda \lambda_i)$$

$$0 < \lambda_i < 1.$$

(See Dyer and Sarin, Keeney and Raiffa, Farquhar, Fishburn or Kirkwood and Sarin for the underlying assumptions and axiomatic derivations of these forms. These authors make a distinction, not covered here, between stochastic and deterministic applications of these results. They usually employ the term value function rather than utility function when referring to deterministic problems).

We would hope that the few scattered instances of application of the explicit utility approach to multiobjective optimization will become more widely noticed and that the approach will receive deeper investigation. There probably will always be some cases when adequate utility assessment is too difficult from a practical viewpoint. But it is probably also fair to say that the multiobjective optimization field has not yet taken sufficient advantage of multiattribute decision theory research. The latter contains a great deal of work in axiomatizing functional forms for utility. It is perhaps a matter of happenstance, not theoretical necessity, that this work has been applied to date mainly in problems with small, enumerable feasible sets.

4.3 Implicit Utility Functions I

Whether or not explicit utility functions are available for multiobjective optimization, some valuable insight can be gained by further

consideration of utility theory. It turns out that some of the multi-objective techniques in current practice are founded on unwittingly specified utility functions. For example, the common procedure of assigning weights w_i and maximizing $\sum_i w_i f_i(x)$, noted in Section 2.1, is nothing more than constructing a linear utility function: $U(F) = \sum_i w_i f_i$. Another example is goal programming, whose underlying utility function is a very rigidly defined piece-wise linear function (Dyer [1977], Rosenthal [1982]).

The choice of a linear utility function has the disturbing implication that the marginal utility of f_i is constant: $\partial U(F)/\partial f_i = w_i$. This conflicts with many generations of economic thought (including Bernoulli's St. Petersburg paradox) which established that these marginal values should not be constant (Baumol). The economic idea is that a decision maker values the next unit of f_i more when f_i is scarce than when it is plentiful. This law of diminishing marginal value can be expressed as $\partial^2 U(F)/\partial f_i^2 < 0$, but with linear utility the second derivatives are of course zero.

The economic inconsistency of applying constant weights to the f_i is perhaps most clearly seen by considering the implications of linear utility on the marginal rates of substitution. If $U(F) = \sum_i w_i f_i$ then $MRS_{ij}(F) = w_i/w_j$, a constant. This implies that the decision maker is indifferent to the amounts of f_i and f_j on hand when considering the rate at which he or she is willing to make tradeoffs between f_i and f_j . This hardly seems reasonable: a unit of commodity is always easier to trade away when the commodity is plentiful than when it is scarce.

Goal programming, the second example given above of a multiobjective technique with an unintended underlying utility function, has fundamentally

the same drawbacks as the constant-weights approach. Goal programming is a very popular multiobjective technique (e.g., Ignizio [1977]) which calls for targets b_i on each f_i and then minimizes a weighted sum of the deviations $d_i(x) = b_i - f_i(x)$. This may be formulated as

$$\min \sum_i [v_i \max(0, d_i(x)) + w_i \max(0, -d_i(x))]$$

s.t. $x \in X$. When $v_i \neq w_i$ this implies a difference in attitude toward under- and over-achievement of the target. It can be easily verified that this model is consistent with the maximization of a utility function for which

$$\frac{\partial U}{\partial f_i}(F(x)) = \begin{cases} v_i & \text{if } d_i(x) < 0 \\ -w_i & \text{if } d_i(x) > 0. \end{cases}$$

Differentiating further, $\partial^2 U(F)/\partial f_i^2 = 0$ (except where it is undefined at $f_i = b_i$). These equations demonstrate that the goal programming model, like linear utility, ignores the normal human tendency to let the amount of f_i on hand influence the marginal value. Similarly, the goal programming framework fails to reflect the fact that a decision maker's willingness to trade off f_i for f_j is usually very dependent on his or her current relative attainments of f_i and f_j . (This critique also applies to pre-emptive goal programming which can be regarded theoretically--though not handled computationally--as a case of the above model in which $w_1 \gg w_2 \gg w_3$ etc. and $v_1 \gg v_2 \gg v_3$ etc. See Dyer [1977] or Rosenthal [1982] for additional comments on goal programming and for computationally oriented observations).

4.4 Implicit Utility Functions II

A number of authors have taken the approach of deliberately rather than unwittingly assuming an implicit utility function. Representative of this class are Geoffrion, Dyer and Feinberg [1972], Harrison [1983], Musselman and Talavage [1980], Oppenheimer [1978], Wehrung [1978], and Zionts and Wallenius [1976, 1983].

As noted earlier, Geoffrion et al. were the first to make a direct attempt at utility maximization without explicitly specifying a utility function. The idea of their approach is to use a nonlinear programming algorithm to attempt solution of

$$\max U(F(x)) \quad \text{s.t. } x \in X$$

in spite of the lack of explicit knowledge of U . (The particular nonlinear programming algorithm chosen was the Frank-Wolfe method, but almost any other primal method such as the reduced gradient algorithm could be used.) When the algorithm calls for information about U , it is obtained through computer/decision maker interaction. This dialogue can be structured so that the decision maker's only task is to answer questions of the form: "which do you prefer, $F(y)$ or $F(z)$, or are you indifferent?" This information is sufficient to enable all other steps of the nonlinear programming algorithm to be executed by the computer (e.g., Dyer [1973]). Remarkably, under the assumption of concave utility (which is consistent with economic theory, e.g. Baumol, this procedure terminates optimally.

The Musselman-Talavage approach also requires concavity and uses interaction to approximate the gradient. It differs in the mathematical programming framework, taking a cutting plane approach to progressively remove inferior points from further consideration. Zionts and Wallenius

have added extra machinery to ensure that an efficient point is achieved and to hasten convergence. Harrison [1983] has also enhanced the original implicit approach, by combining it with ideas related to the explicit utility approach. Oppenheimer's [1978] idea of a proxy function is similar.

In its various forms, the class of methods based on maximizing implicit utility functions appear to be another promising approach for multiobjective optimization. The keystone of the implicit utility approach is the interactively determined gradient, which can be viewed as a correction of the deficiency of the simple weighting (i.e., linear utility) approach. As noted earlier, the weakness of that approach is its ignorance of the fact that an objective has less marginal value, and hence is easier to trade away, when it is plentiful than when it is scarce. The progressively reevaluated gradient is just a temporary, changeable set of weights which account for this fact.

4.5 Relationship Between Efficiency and Utility

Our presentations of efficiency and utility, while taking up the bulk of the paper, have shown very little connection between the two concepts. In terms of the evolution of multiobjective techniques, this separation is probably acceptable because most practical methods are founded on efficiency or utility but not both. It would be incorrect, however, to infer that these two principal ideas of the field are unrelated. Yu [1973, 1974] has developed a theory of cone convexity and domination structures with which he shows that the contrasting approaches of maximizing utility and identifying the efficient set can actually be viewed as (opposite extreme) variations on the same theme. See also Hazen and

Morin [1983] for extension of Yu's results, and Hazen [1983] for consideration of the situation when only partial information about preferences can be obtained.

5. Summary and Conclusions

The numerous multiobjective optimization techniques in existence can be regarded as various ways of synthesizing the principal ideas discussed here. In conclusion, there is certainly no point in strenuously advocating any one method for solving multiobjective problems. Because of the ill definition and variety of the problems and because of the varying abilities of decision makers to articulate preferences, no one method will ever be guaranteed to always find a most preferred solution. The analyst should be prepared to use a number of different techniques. Some general recommendations are given below for his or her assistance:

1. Do not dismiss single-objective optimization. Authors in the multiobjective literature have much too often overstated the case for their field to the extent of virtually writing off single-objective optimization as useless. Besides alienating a group of researchers whose collective efforts have in fact yielded substantial value, they commit the error of closing the door on a potentially viable approach. A carefully designed study based on multiple runs of a single-objective model, with intelligent postoptimality analyses, may sometimes yield as much or more insight as any multiobjective method.
2. Complete generation of the efficient set in the bi-objective case. If there are only two objective functions, the problem is considerably easier than when there are more. The complete

efficient set can be generated parametrically, and, even more important, it can be represented by a simple curve which any manager would find useful. This should be kept in mind when formulating a model, but if a bi-objective model cannot fit the problem, then the following classes of approaches are suggested.

3. Partial generation of the efficient set. Under this rubric, we have unified a broad spectrum of multiobjective techniques which have been applied effectively. At one end of the spectrum are the highly unstructured, totally interactive methods of Soland and Hultz et al. At the other end are the highly structured, totally analytical methods of Ecker and Shoemaker. The interior of the spectrum contains a number of other promising approaches including the methods of Steuer et al., Ho, Gilbert et al. and others.
4. Explicit utility maximization. The fields of multiobjective mathematical programming and multiattribute utility theory have developed simultaneously with hardly any interaction, despite their common aims. Combining the two areas offers a very promising but much overlooked approach. One can use multiattribute decision theory to assess a utility function, and then use a mathematical programming algorithm to maximize utility. Multiattribute utility theory has been used extensively in cases when the feasible region is small enough to be totally enumerated. In the large-feasible-set case, this idea has been used in only the handful of instances listed in Section 4.2, with perhaps very little notice. These few past

experiences and the recent theoretical and empirical results of Harrison demand deeper investigation of this approach.

5. Implicit utility maximization. This is the label we have used for the popular and promising class of methods introduced by Geoffrion, Dyer and Feinberg and significantly extended by Zionts and Wallenius, Oppenheimer, Harrison and others. (It was also noted that some other techniques, like simple weighting and goal programming, have underlying implicit utility assumptions, which should be questioned carefully by the analyst before being embraced.)

It would be nice to be able to close with a simple algorithm for advising the prospective problem solver on which of these recommended classes of approaches to pursue first, but, unfortunately, life is never that easy.¹

¹Acknowledgements. This paper has benefited from the comments of Les Foulds, Terry Harrison, Benjamin Hobbs and Timothy Lowe. I would also like to thank Candance Wages and her staff for their careful production of the paper.

References

1. Baumol, W. J., [1977], Economic Theory and Operations Analysis, 4th ed., Prentice-Hall.
2. Benson, H. P. [1978], "Existence of Efficient Solutions for Vector Maximization Problems," Journal of Optimization Theory and Applications 26, 569-580.
3. Benson, H. P., [1979], "Vector Maximization with Two Objective Functions," Journal of Optimization Theory and Applications 28, 253-257.
4. Benson, H. P. and T. L. Morin, [1977], "The Vector Maximization Problem: Proper Efficiency and Stability," SIAM Journal on Applied Mathematics 32, 64-72.
5. Cochrane, J. L. and M. Zeleny (eds.) [1973], Multiple Criteria Decision Making, University of South Carolina Press, Columbia, S.C.
6. Cohon, J. L., [1978], Multiobjective Programming and Planning, Academic Press.
7. Cohon, J. L. and D. M. Marks, [1975], "A Review and Evaluation of Multiobjective Programming Techniques," Water Resources Research 11, 208-220.
8. Daellenbach, H. G., [1981], "Multiple Objective Decision Making," Operations Research Department, University of Canterbury, Christchurch, New Zealand.
9. DeWispelare, A. R. and A. P. Sage, [1981], "On Combined Multiple Objective Optimization Theory and Multiple Attribute Utility Theory for Evaluation and Choicemaking," Large Scale Systems 2, 1-19.
10. Dyer, J. S., [1973], "An Empirical Investigation of a Man-Machine Interactive Approach to the Solution of the Multiple Criteria Problem in Cochrane and Zeleny, op. cit., 202-216.
11. Dyer, J. S., [1977], "On the Relationship Between Goal Programming and Multiattribute Utility Theory," Discussion Paper No. 69, Management Science Study Center, Graduate School of Management, U.C.L.A.
12. Dyer, J. S. and R. K. Sarin [1979], "Measurable Multiattribute Value Functions," Operations Research 27, 810-822.
13. Ecker, J. G. and I. A. Kouada, [1975], "Finding Efficient Points for Multiple Objective Linear Programs," Mathematical Programming 8, 375-377.

14. Ecker, J. G. and I. A. Kouada, [1978], "Finding All Efficient Extreme Points for Multiple Objective Linear Programs," Mathematical Programming 14, 249-261.
15. Ecker, J. G., N. S. Hegner and I. A. Kouada, [1980], "Generating All Maximal Efficient Faces for Multiple Objective Linear Programs," Journal of Optimization Theory and Applications 30, 353-381.
16. Ecker, J. G. and N. E. Shoemaker [1981], "Selecting Subsets from the Set of Nondominated Vectors in Multiple Objective Linear Programming," SIAM Journal of Control and Optimization 19, 505-515.
17. Evans, G. W., [1983], "An Overview of Techniques for Solving Multi-objective Mathematical Programs," Working Paper, Industrial Engineering Department, University of Louisville.
18. Farquhar, P. H., [1983], "Research Direction in Multiattribute Utility Analysis," in Essays and Surveys on Multiple Criteria Decision Making, Pierre Henson, editor. Lecture notes in Economics and Mathematical Systems, No. 209, Springer Verlag, Berlin.
19. Farquhar, P. H., [1984], "Utility Assessment Methods," to appear in Management Science, November 1984.
20. Fishburn, P. C., [1983], The Foundations of Expected Utility, Reidel, Dordrecht, The Netherlands.
21. Gal, T., [1977], "A General Method for Determining the Set of All Efficient Solutions to a Linear Vector Maximum Problem," European Journal of Operations Research, 307-322.
22. Geoffrion, A. M., [1968], "Proper Efficiency and the Theory of Vector Maximization," Journal of Mathematical Analysis and Applications 22, 618-630.
23. Geoffrion, A. M., J. S. Dyer and A. Feinberg, [1972], "An Interactive Approach for Multicriterion Optimization with an Application to the Operation of an Academic Department," Management Science 19, 357-368.
24. Gilbert, K., D. D. Holmes and R. E. Rosenthal, [1982], "A Multi-objective Discrete Optimization Model for Land Allocation" Working Paper No. 155, Management Science Program, The University of Tennessee.
25. Gilbert K. C. and R. M. Shane, [1981], "TVA Hydro Scheduling Model: Theoretical Aspects," Journal of Water Resources Planning and Management Division, American Society of Civil Engineers, Vol. 108, No. WR1, 1-19.
26. Golabi, K., C. W. Kirkwood and A. Sicherman [1981], "Selecting a Portfolio of Solar Energy Projects Using Multiattribute Preference Theory," Management Science 27, 174-189.

27. Goicoechea, A., D. Hansen and L. Duckstein, [1982], Multiobjective Decision Analysis with Engineering and Business Applications, John Wiley and Sons, New York.
28. Gros, J., [1975], "Power Plant Siting: A Paretian Environmental Approach," Nuclear Engineering and Design 34, 281-292.
29. Haimes, Y., D. Hall and S. Freedman, [1975], Multiobjective Optimization in Water Resources Systems, Elsevier, Amsterdam.
30. Hannan, E. L., J. A. Smith and G. R. Gilbert, [1983], "A Multi-attribute Decision Making Approach to the Selection of an Auxilliary Device for Icebreakers," Decision Sciences 14, 240-252.
31. Harrison, T. P., [1983], "Design and Development of a Decision Support System for Forest Management Planning with Multiple Objectives," Ph.D. Dissertation, Management Science Program, The University of Tennessee.
32. Harrison, T. P. and R. E. Rosenthal, [1981], "A Multiobjective Optimization Program for Scheduling Timber Harvests on Nonindustrial Private Forest Lands," Working Paper No. 119, Management Science Program, The University of Tennessee. (To appear in TIMS Study on Systems Analysis in Forestry and Forest Industries.)
33. Harrison, T. P. and R. E. Rosenthal, [1984], "On the Use of Utility/Value Functions in a Multiobjective Mathematical Programming Framework," Working Paper No. 84-6, Management Science Department, Pennsylvania State University.
34. Hazen, G. B., [1983], "Partial Information, Dominance, and Potential Optimality in Multiattribute Utility Theory," Working Paper, Northwestern University, Evanston, Illinois.
35. Hazen, G. B. and T. L. Morin, [1983], "Optimality Conditions in Nonconical Multiple-Objective Programming," Journal of Optimization Theory and Applications 40, 25-60.
36. Ho, J. K., [1979], "Holistic Preference Evaluation in Multiple Criteria Optimization," available through Management Science Program, The University of Tennessee.
37. Ho, J. K., [1980], "Multiple Criteria Optimization: A Unified Framework," available through Management Science Program, The University of Tennessee.
38. Hobbs, B. F., [1979], "Analytical Multiobjective Decision Methods for Power Plant Siting: A Review of Theory and Applications," Report BNL-NUREG-51204, Brookhaven National Laboratory.
39. Hobbs, B. F., [1980], "A Comparison of Weighting Methods in Power Plant Siting," Decision Sciences 11, 725-737.

40. Hultz, J. W., D. D. Klingman, G. T. Ross and R. M. Soland, [1981], "An Interactive Computer System for Multicriteria Facility Location," Computers and Operations Research 8, 249-261.
41. Hwang, C. L., A. S. M. Masud, S. R. Paidy and K. Yoon, [1979], Multiple Objective Decision Making -- Methods and Applications, Springer-Verlag.
42. Ignizio, J., [1977], "A Review of Goal Programming: A Tool for Multiobjective Analysis," Journal of the Operational Research Society 29, 1109-1119.
43. Isermann, H., [1974], "Proper Efficiency and the Linear Vector Maximum Problem," Operations Research 22, 189-191.
44. Keefer, D. L., [1978], "Applying Multiobjective Decision Analysis to Resource Allocation Planning Problems," in Zions op. cit., 299-320.
45. Keeney, R. L., [1977], "The Art of Assessing Multiattribute Utility Functions," Organizational Behavior and Human Performance 19, 267-310.
46. Keeney, R. L., [1979], "Evaluation of Proposed Pumped Storage Sites," Operations Research 27, 48-64.
47. Keeney, R. L., [1980], Siting Energy Facilities, Academic Press, New York.
48. Keeney, R. L. and H. Raiffa, [1976], Decisions with Multiple Objectives: Preferences and Value Tradeoffs, John Wiley and Sons.
49. Kirkwood, C. W. and R. K. Sarin, [1980], "Preference Conditions for Multiattribute Value Functions," Operations Research 28, 225-232.
50. Kuhn, H. W. and A. W. Tucker, [1950], "Nonlinear Programming," Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability (J. Neyman, ed.), 481-492.
51. Musselman, K. and J. Talavage, [1980], "A Tradeoff Cut Approach to Multiple Objective Optimization," Operations Research 28, 1424-1435.
52. Oppenheimer, K. R., [1978], "A Proxy Approach to Multi-Attribute Decision Making," Management Science 24, 675-689.
53. Ringuest, J. L. and T. R. Gullledge, [1983], "A Preemptive Value Function Method Approach for Multiobjective Linear Programming Problems," Decision Sciences 14, 76-86.
54. Rosenthal, R. E., [1983], "Goal Programming: A Critique," New Zealand Operational Research, 11, 1-8.

55. Roy, B. and P. Vincke, [1981], "Multicriteria Analysis: Survey and New Directions," European Journal of Operational Research 8, 207-218.
56. Schoemaker, P. J. H. and C. C. Waid [1982], "An Experimental Comparison of Different Approaches to Determining Weights in Additive Utility Models," Management Science 28, 182-196.
57. Simon, H. A., [1957], Models of Man, John Wiley and Sons, New York.
58. Simon, H. A., [1964], "On the Concept of Organizational Goal," Administrative Science Quarterly 9, 1-22.
59. Soland, R. M., [1979], "Multicriteria Optimization: A General Characterization of Efficient Solutions," Decision Sciences 10, 26-38.
60. Steuer, R. E., [1976], "Multiple Objective Linear Programming with Interval Criterion Weights," Management Science 23, 305-317.
61. Steuer, R. E. and E. U. Choo, [1983], "An Interactive Weighted Tchebycheff Procedure for Multiple Objective Programming," Mathematical Programming 26, 326-344.
62. Steuer, R. E. and A. T. Schuler, [1978], "An Interactive Multiple-Objective Linear Programming Approach to a Problem in Forest Management," Operations Research 26, 254-269.
63. Wehrung, D. A., [1978], "Interactive Identification and Optimization Using a Binary Preference Relation," Operations Research 26, 322-331.
64. White, D. J., [1980], "Multi-Objective Interactive Programming," Journal of the Operational Research Society 31, 517-523.
65. Winkler, R. L., [1982], "Research Directions in Decision Making Under Uncertainty," Decision Sciences 13, 517-533.
66. Yu, P. L., [1973], "Introduction to Domination Structure in Multi-criteria Decision Problems," in Cochrane and Zeleny, op. cit., 249-261.
67. Yu, P. L., [1974], "Cone Convexity, Cone Extreme Points and Non-dominated Solutions in Decision Problems with Multiobjectives," Journal of Optimization Theory and Applications 14, 319-377.
68. Yu, P. L. and M. Zeleny, [1975], "The Set of All Non-Dominated Solutions in Linear Cases and a Multicriteria Simplex Method," Journal of Mathematical Analysis and Applications 49, 430-468.
69. Zeleny M., [1982], Multiple Criterion Decision Making, Prentice-Hall. (This book contains a very extensive bibliography.)
70. Zionts, S. (ed.), [1978], Multiple Criteria Problem Solving: Proceedings, Buffalo, N.Y. (U.S.A.), 1977, Springer-Verlag, New York.

71. Zionts, S. and J. Wallenius, [1976], "An Interactive Programming Method for Solving the Multiple Criteria Problem," Management Science 22, 652-663.
72. Zionts, S. and J. Wallenius, [1983], "An Interactive Multiple Objective Linear Programming Method for a Class of Underlying Nonlinear Utility Functions," Management Science 29, 519-529.

DISTRIBUTION LIST

	NO. OF COPIES
Defense Technical Information Center Cameron Station Alexandria, VA 22314	2
Library Code 0142 Naval Postgraduate School Monterey, CA 93943	2
Research Administration Code 012A Naval Postgraduate School Monterey, CA 93943	1
Library Code 55 Naval Postgraduate School Monterey, CA 93943	1
Professor Richard E. Rosenthal Code 55 Naval Postgraduate School Monterey, CA 93943	50
Professor Gerald Brown Code 55 Naval Postgraduate School Monterey, CA 93943	1
Dr. Roland W. Kinney National Research Council National Academy of Sciences 2101 Constitution Ave, NW Washington, D.C. 20418	1

END

FILMED

12-30-41

DTIC